

THE CLAIMS

What is claimed is:

1. A method for asynchronously remotely copying database content changes from a primary site to a remote site, the method comprising:

associating a sequential identification with each respective log record write and each corresponding data record write received at the primary site;

asynchronously remotely copying each respective log record write from the primary site to the remote site;

receiving an acknowledgement at the primary site, the acknowledgement corresponding to a log record write that has been completed at the remote site; and

asynchronously remotely copying each data record write having a sequential identification that is prior to or equal to the sequential identification of the log record write corresponding to the received acknowledgement.

2. The method according to claim 1, wherein the sequential identification is a monotonically increasing identification number.

3. The method according to claim 1, wherein the sequential identification is a monotonically increasing time-stamp identification.

4. The method according to claim 1, wherein a log record write is asynchronously remotely copied from the primary site to the remote site before a data record write is asynchronously remotely copied from the primary site to the remote site.

5. The method according to claim 1, wherein each log record write is a log block and each data record write is a data block write.

6. The method according to claim 1, further comprising:

asynchronously receiving a log record write at the remote site;
storing the received log record write at the remote site;
sending an acknowledgement from the remote site to the primary site when the received log record write is complete;

asynchronously receiving a data record write at the remote site from the primary site;
and

storing the received data record write.

7. A method for asynchronously remotely coping database content changes occurring at a primary site at a remote site, the method comprising:

asynchronously receiving a log record write at the remote site, each respective log record received at the remote site having an associated sequential identification and a corresponding data record write;

storing the received log record write at the remote site;
sending an acknowledgement from the remote site to the primary site when the received log record write is complete;

asynchronously receiving a data record write at the remote site from the primary site, each received data record write having a sequential identification that is prior to or equal to the sequential identification of the log record write corresponding to the received acknowledgement; and
storing the received data record write.

8. The method according to claim 7, wherein the sequential identification is a monotonically increasing identification number.

9. The method according to claim 7, wherein the sequential identification is a monotonically increasing time-stamp identification.

10. A storage system for asynchronously remotely copying content changes stored in the

storage system, the system comprising:

a primary site having a storage system separately storing log records and data records;
a remote site having a storage system separately storing log records and a data records,

the primary site associating a sequential identification with each respective log record write and each corresponding data record write occurring at the primary site and asynchronously remotely copying each respective log record write from the primary site to the remote site, the remote site sending to the primary site an acknowledgement corresponding to a log record write that has been completed at the remote site, and the primary site asynchronously remotely copying to the remote site each data record write having a sequential identification that is prior to or equal to the sequential identification of the log record write corresponding to the received acknowledgement.

11. The system according to claim 10, wherein the sequential identification is a monotonically increasing identification number.

12. The system according to claim 10, wherein the sequential identification is a monotonically increasing time-stamp identification.

13. The system according to claim 10, wherein a log record write is asynchronously remotely copied from the primary site to the remote site before a data record write is asynchronously remotely copied from the primary site to the remote site.

14. The method according to claim 10, wherein each log record write is a log block and each data record write is a data block write.

15. A primary site of a distributed storage system, the system comprising:
a storage system separately storing log records and data records; and
a controller associating a sequential identification with each respective log record write

and each corresponding data record write occurring at the primary site and asynchronously remotely copying each respective log record write from the primary site to a remote site, the controller receiving an acknowledgement corresponding to a log record write that has been completed at the remote site and, in response, asynchronously remotely copying to the remote site each data record write having a sequential identification that is prior to or equal to the sequential identification of the log record write corresponding to the received acknowledgement.

16. The system according to claim 15, wherein the sequential identification is a monotonically increasing identification number.

17. The system according to claim 15, wherein the sequential identification is a monotonically increasing time-stamp identification.

18. The system according to claim 15, wherein a log record write is asynchronously remotely copied from the primary site to the remote site before a data record write is asynchronously remotely copied from the primary site to the remote site.

19. The method according to claim 15, wherein each log record write is a log block and each data record write is a data block write.

20. A remote site of a distributed storage system, the system comprising:
a storage system separately storing log records and data records,
a controller asynchronously receiving a log record write from a primary site, each respective log record received at the remote site having an associated sequential identification and a corresponding data record write, storing the received log record write in the storage system and sending an acknowledgement from the remote site to the primary site when the received log record write is complete.

21. The remote site according to claim 20, wherein the controller further asynchronously receives a data record write from the primary site, each received data record write having a sequential identification that is prior to or equal to the sequential identification of the log record write corresponding to the received acknowledgement, and stores the received data record write.

22. The remote site according to claim 20, wherein the sequential identification is a monotonically increasing identification number.

23. The remote site according to claim 20, wherein the sequential identification is a monotonically increasing time-stamp identification.

24. The remote site according to claim 20, wherein a log record write is asynchronously remotely copied from the primary site to the remote site before a data record write is asynchronously remotely copied from the primary site to the remote site.

26. The remote site according to claim 20, wherein each log record write is a log block and each data record write is a data block write.